



[12] 发明专利申请公开说明书

[21] 申请号 94194846.3

[43]公开日 1997 年 1 月 29 日

[11] 公开号 CN 1141701A

[22]申请日 94.11.29

[30]优先权

[32]93.11.30[33]FI[31]935365

[86]国际申请 PCT/FI94/00535 94.11.29

[87]国际公布 WO95/15637 英 95.6.8

[85]进入国家阶段日期 96.7.16

[71]申请人 诺基亚电信公司

地址 芬兰埃斯波

[72]发明人 朱哈·帕贾弗塔 海克·萨洛沃奥瑞

乔玛·马卡塞尔卡 米卡·卡斯林

塞波·皮哈拉米 米科·奥科恩

里查德·菲尔曼

米科·莱霍

[74]专利代理机构 中国国际贸易促进委员会专利商标
事务所


代理人 鄢 迅

权利要求书 2 页 说明书 10 页 附图页数 3 页

[54]发明名称 帧中继网络中过载情况的控制

[57]摘要

本发明涉及帧中继网络中的拥挤管理方法和系统。该方法包括当在网络节点收到待发送的帧时确定与该帧相关联的虚拟通道。为了在所有用户之间公正地划分网络的中继容量,(a)在用户节点的输入边界处将数据缓冲到虚拟通道专用的缓冲器(61₁...62_n),(b)在向后方向上从网络节点向其帧在特定时间接收的虚拟通道的用户节点发送拥挤通知(M),以及(c)在具有预定长度的某时段期间从与所述虚拟通道相应的用户节点向网络提供的通信量是根据在所述间隔期间从网络接收的拥挤通知(M)的内容受到控制的。



权 利 要 求 书

1. 帧中继网络中的拥挤管理方法, 该帧中继网络包括用户节点, 用户(A...E)经过数据链路(14a...14e)连接到用户节点上, 所述方法包括当在网络节点接收将要发送的帧(39)时确定与该帧相关联的虚拟通道, 其特征在于:

-在用户节点的输入边界将数据缓冲到虚拟通道专用的缓冲器(62₁...62_n),

-在向后方向将拥挤通知(M)从网络节点发送到其帧在特定时刻接收的虚拟通道的用户节点, 以及

-在预定长度的某间隔期间从与所述虚拟通道相应的用户节点缓冲器向网络提供的通信量根据在所述间隔期间从网络收到的拥挤通知(M)的内容受到控制。

2. 根据权利要求1的方法, 其特征在于拥挤通知的严重度与节点处的缓冲器的填充率有关, 从而某个严重度与每个填充率范围相应, 在所述间隔期间向网络提供的通信量的恒定幅度的改变与拥挤的某个严重度相应。

3. 根据权利要求1的方法, 其特征在于如果在所述间隔期间没有发送属于该虚拟通道的拥挤通知, 则在用户节点处的通信量增加某个恒定值, 然而并不超出所允许的最大值。

4. 根据权利要求3的方法, 其特征在于当通信量超出所提交的突发串尺寸(Bc)的量时, 在收到拥挤通知(M)后, 立即将通信量降至一个与所提交的突发串尺寸(Bc)相应的值。

5. 根据权利要求1的方法, 其特征在于通信量是根据在所述间隔期间所收到的最严重拥挤通知控制的。

6. 根据权利要求1的方法, 其特征在于在发送一个拥挤通知之后, 在一个预定时间内避免向同一虚拟通道发送随后的拥挤通知, 以避免在虚拟通道上产生一个突发串时徒劳地发送几个拥挤通知。

7. 帧中继网络中的拥挤管理系统, 该帧中继网络包括多个在其间

传输数据的节点(N; N1...N4), 至少部分所述节点是用户节点, 网络的用户(A...E)经数据链路(14a...14e)连接到用户节点上, 其特征在于:

- 在用户节点的输入边界设置虚拟通道专用的缓冲器(62₁...62_n), 数据被缓冲到其中,

- 节点包括用于在特定时间在向后方向上向其帧被接收的虚拟通道的用户节点发送拥挤通知(M)的装置, 以及

- 用户节点包括用于在具有预定长度的某间隔期间对从与所述虚拟通道相应的缓冲器向网络提供的通信量进行控制的装置(65,S3), 所述控制是根据在所述间隔期间从网络接收的拥挤通知(M)的内容进行的。

说明书

帧中继网络中过载情况的控制

本发明涉及在帧中继网络中用于拥挤管理的根据权利要求 1 前序部分的方法和根据权利要求 7 前序部分的系统。

拥挤是指这样一种情况：在特定时间，在某网络点(称为瓶颈资源)处传输请求数超过传输容量。拥挤通常导致过载状态，其结果是，例如缓冲器溢出，因此网络或者用户将重新传输包。拥挤管理(CM)的功能是维持传输请求与传输容量之间的平衡，使得瓶颈资源在最优级上工作，并且以确保公正的方式向用户提供服务。

拥挤管理可以分为拥挤避免(CA)和拥挤恢复(CR)。拥挤避免方法的目的在于：通过根据网络拥挤状态动态调节用户的带宽，并且/或者通过向网络路由报警，使得瓶颈资源的部分通信负载转移到空闲资源上，来防止网络中产生拥挤。而恢复方法的目的是：如果避免方法已经不能防止产生拥挤，则将瓶颈资源的工作恢复到最优级。

帧中继(frame relay FR)技术是一种取代当前所用的包交换网络连接用于传输变长帧的包交换网络技术。当前包交换网络中广泛采用的协议(X.25)需要足够的处理，并且传输设备昂贵，其结果是速度也低。这些情况是因为这一事实：X.25 标准是在所用的传输连接仍然非常容易出现传输错误时开发的。帧中继技术的出发点是相当低的传输线路错误概率。因此已经能够放弃帧中继技术中的某些不必要的功能，这使帧传递迅速且有效。帧方式载体服务(Frame Mode Bearer Service)一般在 CCITT 荐议 I.233(参考文献 1)和荐议 Q.922(参考文献 2)中相关的协议中描述。FR 网络中的拥挤和拥挤管理机制在 CCITT 荐议 I.370(参考文献 3)中描述。为了更详细地描述 FR 技术，请参考 1991 年 4 月 McGrawHill 公司的 Datapro Management of Data Communications 的 An Overview of Frame Relay Technology(参考文献 4)以及上

述荐议。

在当前所用的 FR 网络分级结构中，节点具有仅与物理通道相应的接收和发送缓冲器，即，几个不同虚拟通道的通信和应用经过相同的缓冲器。对来自节点的链路输出，尽可能地排空缓冲器，而不考虑网络中总的拥挤程度。因此，尽可能多地加载来自节点的链路输出，即使在较接近网络中心(在拥挤节点处)时将可能放弃该帧。除了浪费网络资源之外，在网络中继节点处放弃帧影响以较长吞吐量延迟的形式利用网络的应用系统(虚拟通道是指具有一条传输链路长度的虚拟连接部分，而虚拟连接是实际的包交换端对端 FR 连接)。

即使在用户连接处根据通用服务参数 B_c 、 B_e 和 CIR 对自虚拟通道所接收的通信进行监测，当从用户节点传送帧时，该监测也不起作用。[参数 B_c (提交的脉冲串大小)代表在某时隙 T_c (一般 $T_c=1$ 秒)内在网络上用户能够传输的最大数据量；参数 B_e (剩余脉冲串大小)代表在时隙 T_c 内用户能够超过值 B_c 的数据量；而参数 CIR(提交的信息率)代表网络在正常状态下所保证的数据传输率， $CIR=B_c/T_c$ 。这些参数在参考文献 3 中定义。] 帧被从接收缓冲器按选定路线送到正确的通道专用的传输缓冲器。到达节点的帧从而根据 FIFO 原则经过整个节点。由于帧中继网络中通信的突发性，从上述经常可以推定一条虚拟通道占用一个节点的缓冲和中继容量的主要部分。这影响了向该节点提供通信的具有更长吞吐量延迟和更高帧丢失概率的其他虚拟通道。甚至一个具有高突发性或不可满足性的虚拟通道可能使该连接的其他虚拟通道的服务水平下降很大。

本发明的目的是避免上述缺陷，并提供 FR 网络中的一种新型的拥挤管理方法，以允许比以前更公正地和更有效地利用网络资源。这是通过本发明的方法实现的，其特征在于权利要求 1 的特征部分中所公开的内容。本发明的系统的特征在于权利要求 7 的特征部分中所公开的内容。

本发明的构思是，将拥挤通知从网络节点传输到其帧在该节点

接收的虚拟通道的用户节点,并例如根据这些拥挤通知所指示的网络中的拥挤等级,在用户节点的虚拟通道专用缓冲器中控制由每条虚拟通道从用户节点向网络提供的通信量。

本发明的方法允许在所有用户之间公正划分单个节点,以及最终整个网络的中继容量。

下面,将参照附图中所示的例子更详细地描述本发明及其最佳实施方式,其中

图1示出根据本发明的方法的典型工作环境,

图2示出根据本发明的FR网络用户节点,

图3示出FR网络中待传送的帧的格式,

图4示出网络中拥挤通知的传递,

图5a示出在第一示例情形下根据拥挤通知改变带宽,

图5b示出在第二示例情形下根据拥挤通知改变带宽,以及

图6示出根据本发明的FR网络中继节点。

帧中继网络能够由几种不同的应用系统使用,它们不需要类似的服务。因此,考虑到两个最重要的参数(帧丢失概率和延迟),在根据应用系统将服务分为不同种类的网络中,采用本发明的方法是有利的。这样一种解决方案公开在芬兰专利申请第925671号中。在该申请中,提出将服务分为以下三类:

- 第一服务类(种类1)提供交互式服务,延迟短,
- 第二服务类(种类2)提供低的帧丢失概率,没有任何明显察觉的延迟,
- 第三服务类(种类3)既提供短的延迟又提供低的帧丢失概率。

以这种方式实现的网络的每个中继连接具有服务种类专用的缓冲器,每个服务种类一个。而一个用户节点在用户连接侧具有虚拟通道专用的缓冲器。下面将更详细地描述这些解决方案;另外请参照以上所引用的芬兰专利申请。

图1示出提供公众网络服务的FR网络,即对单个公司或多个公司的不同办公室A...E的局域网络11进行互连的帧中继网络

12。每个办公室的局域网络 11 经过一个局域网桥 13 和分别由标号 14a...14e 所指的数据链路使用 FR 服务。在 FR 用户 A...E 与 FR 网络节点 N 之间的连接自身是公知的，因此在此不做更详细地描述。有关在互连中所用的局域网络和桥的更详细的信息可见如 1991 年 2 月《电信》中 Michael Grimshaw LAN Interconnections Technology 的文章，并且见 1991 年的 Lähiverkko-opas, Leena Jaakonmäki, Suomen ATK-kustannus Oy，在此将它们作为参考文献。

FR 网络的已知节点结构的一般特征是对所有帧使用相同缓冲器，假定它们被按选定路线送到相同的物理连接。相反，根据本发明，在所有网络节点的输出边界和具有中继连接的输入边界设置与上述服务种类相应的缓冲器。图 2 示出在网络中一个中继节点处的这类解决方案。该节点接收在用户连接的桥 13(图 1)中所原始装配的 FR 帧。在桥 13 中，用户 LAN11 的帧插到 FR 帧的信息域中(时序位和其他类似位例外)。图 3 示出在 FR 帧 39 的信息域中插入 LAN 帧 38。也示出一种一般 FR 网络帧格式，在信息域之前的地址域包括两个八位字节(位 1 至 8)。第一个八位字节的位 3 至 8 和第二个八位字节的位 5 至 8 构成一个数据链路连接标识符 DLCI，它向节点指示例如一个特定帧所属的虚拟连接和虚拟通道。虚拟通道通过数据链路连接标识符相互区分。然而，数据链路连接标识符仅在单个虚拟通道上是非歧义性的，在到下一虚拟通道的转变中，它在节点中可能改变。第二地址域八位字节的第 2 位，称为 DE 位(放弃合格指示符)，对于帧的放弃也是重要的。根据 CCITT 荐议，例如在拥挤状态下，如果一个帧的 DE 位已被置为 1，则允许放弃该帧。因为 FR 帧中的其他位与本发明无关，所以在此不对它们做更详细的说明。为了更详细地描述，参考上述参考文献 2 和 4。

在网络外围的用户节点(图 2)处，用户连接 14a、14b 等(在图 2 所示的例子中，它们连接到同一节点)，首先连接到一个标识部件 61，该部件接收在桥 13(图 1)中所形成的 FR 帧。该标识部件

61 从帧的地址域中读数据链路连接标识符 DLCI, 并将该帧送到与标识符所指示的虚拟连接相应的输入缓冲器 $62_1 \dots 62_n$ 。每条数据链路具有一个专用选择器 S3, 它从每个虚拟通道的输入缓冲器中选择帧, 并将帧送到集中式路由器 16, 该路由器又将帧送到正确数据链路(图中只示出一条输出数据链路)的分类部件 43。分类部件 43 从帧的地址域中读标识符 DLCI, 并从表 T 中选择与标识符所指的虚拟通道相应的服务种类。在分类已经完成的基础上, 分类部件 43 将每个帧加到与该帧的服务种类相应的输出缓冲器 64a、64b 或 64c。每条输出数据链路因此具有三个输出缓冲器, 每个服务种类一个。选择器 S4 从这些服务种类专用的缓冲器选择帧, 并将帧送到中继连接。

由用户在 FR 网络上传输的通信如此在每条虚拟连接专用的用户节点的输入侧得以缓冲。输入帧 39 在每条虚拟连接上动态地链接起来。根据虚拟连接的服务种类, 链长度具有预定的可允许最大值; 在服务种类 1 和 3 时该值较小, 而在服务种类 2 时该值较大。选择器 S3 读例如与分配给它们的通信量成比例的缓冲器 $62_1 \dots 62_n$ 等, 从而满足公正原则。

根据本发明, 由每条虚拟通道提供到网络 12 的通信量在虚拟通道专用的缓冲器 $62_1 \dots 62_n$ 中受到调节, 使得通信量根据网络中拥挤的等级围绕分配给该通道的基值而变化。因此, 提供给网络内部的通信总量是以分配给虚拟通道的服务等级、带宽和网络中总的拥挤等级为基础的, 其中, 服务等级和带宽确定某个基值, 围绕该值, 根据网络中总的拥挤等级调节通信量。该调节是通过从网络向其帧被收到的虚拟通道的用户节点传输拥挤通知来执行的; 在预定长度的给定时间间隔内, 从与该虚拟通道相应的用户节点缓冲器向网络提供的通信量, 根据在该间隔内从网络收到的拥挤通知的内容受到控制。图 2 中由标号 M 所指的拥挤通知连接到用户节点中控制选择器 S3 的控制部件 65, 所述控制部件根据拥挤通知的内容控制从每条虚拟通道专用的缓冲器所读的数据的量。每个缓冲器因而根据其自身参数和其自身拥挤通知而排空。

对于用户节点，本发明的方法能够比为闸门：每条虚拟通道专用的缓冲器允许将向网络发送的通信，由起闸门作用的系统控制（通过控制部件 65），该系统限制从缓冲器读的数据量，即调节虚拟通道的带宽。

该闸门是与上述服务参数 B_c 相联的系统，在某个时间周期内，它只允许从虚拟通道专用的缓冲器向网络传输某些量的数据送到网络。至于监测虚拟通道带宽的方案，能够使用具有例如 $T_c/3$ (T_c 具有经常为 1 秒的长度) 长度的时间间隔。在具有 $T_c/3$ 长度的每个间隔期间，将缓冲器向朝向网络的输出连接排空，其量为闸门高度所允许的量。剩余的帧留在缓冲器中，等待下一 $T_c/3$ 的间隔。

根据本发明，网络的节点在向后的方向上将上述拥挤通知 M 传输给用户节点。这些拥挤通知指示在该节点处缓冲器的填充率。这种机制示于图 4。网络的节点由标号 $N1...N4$ 指示，用户由标号 S 指示。当在网络节点（例如图 2 中的节点 $N1$ ）服务种类专用的缓冲器超过某填充率，例如 20%，则该节点在向后的方向上将一个拥挤通知传输给其帧被收到的虚拟通道的用户节点（图 4 中的 $N2$ ）。节点 ($N1$) 如上所述利用数据链路连接标识符 $DLCI$ 标识正确的虚拟通道。

在用户节点确定每条虚拟通道专用缓冲器的排空率的闸门高度（即提供给虚拟通道的带宽）根据在所述虚拟通道上从网络接收的拥挤通知的内容以这样一种方式得到调节，使得较高的负载对应较小的带宽。在间隔 $T_c/3$ 期间对带宽的调节最好是根据在所述间隔期间所收到的最严重的拥挤通知执行的。如果虚拟通道在间隔 $T_c/3$ 期间内未收到任何拥挤通知，则根据本发明，在下一间隔开始时，自动增大该通道的带度。该增大可以确定为例如 10%。带宽的初值是为该通道分配的值 $B_c + B_e$ ，即在任何情况下都不允许超出的最大通信量。

通过拥挤通知，能够指示例如三种不同等级的拥挤，从而与它们相应的带宽的改变（都根据最大值计算），例如可以与下表中所指的类似。

拥挤通知的严重程度	带宽的改变
1	-10%
2	-15%
3	-20%

为说明上述原理，图 5a 和 5b 示出根据本发明的调节的两个不同例子。纵轴代表闸门高度(即带宽)，而横轴代表时间，由长度为 $T_c/3$ 的连续间隔组成。在图 5a 所示的情形下，首先收到严重度为 1 的拥挤通知 M，从而立即降低带度。之后收到严重度为 2 的第二个拥挤通知 M，从而再次立即降低带度。最后，在间隔 $T_c/3$ 期间收到具有严重度为 3 的第三个拥挤通知 M，从而又立即降低带宽(于是具有最低可能值)。因为在随后的间隔 $T_c/3$ 期间未收到拥挤通知，所以在该间隔之后(在随后间隔 $T_c/3$ 的开始处)，通道的带宽自动增大(时刻 T2)。

在图 5b 所示的情形下，在时域内拥挤通知的次序相反。首先收到严重度为 3 的拥挤通知，从而立即将带宽降到最低可能值。在其他间隔期间收到的拥挤通知不再改变带宽，而不管其严重度如何。如果首先收到严重度为 2 的拥挤通知，则只有严重度为 3 的拥挤通知造成进一步降低带宽。

根据本发明，当在当前间隔期间内收到目前最严重的拥挤通知后，立即降低带宽，但是只有在间隔 $T_c/3$ 期间未收到拥挤通知时，带宽才会增加。在该情形下，在当前间隔已经结束后开始增加。

如果带宽已经超出值 B_c ，则在收到第一个拥挤通知后，根据本发明的一个最佳实施方式，它立即降至值 B_c ，因而在该特定情形下，不遵循为改变设定的上述限制。这保证最终对瞬时拥挤进行快速反应。然而，在所有情形下，带宽是如上所述向上调节的。

与位于网络节点的缓冲器的填充率相关的拥挤通知的严重性的每个程度，可以设定为例如与下表所给的填充率限制相应：

总缓冲容量上所计算的填充率	拥挤严重度
20...39%	1
40...59%	2
60...100%	3

网络的每个节点连续监测服务种类专用缓冲器的填充率。当在拥挤节点收到一个新帧时，该节点在该帧接收方向上向当前虚拟通道的用户节点发送一个拥挤通知 M。如果缓冲器的填充率小于 20%，则输入帧并不导致发送拥挤通知。在发送拥挤通知的同时，在该节点由一个定时器设置预定长度的间隔 T1；在该间隔期间，不向所涉及的虚拟通道发送新的拥挤通知。间隔 T1 的长度例如可以是 100ms(即大约是间隔的三分之一 $T_c/3$)。以这种方式，确保在该虚拟通道上产生突发串时不徒劳地发送几个拥挤通知。当定时器计时结束时，如果必要，则可能再次向虚拟通道发送拥挤通知。

拥挤通知 M 应极快地传送到出口节点，以使能够对拥挤做出快速反应。根据本发明最佳实施方式，这些拥挤通知形成一个单独的第四服务种类，为此在节点处设置单独的服务种类专用的缓冲器。对于用户节点，这意味着节点输出侧设置有第四输出缓冲器，在图 2 中标以虚线和标号 64d。

对于网络的中继节点，该实施方式标示图 5，其中与拥挤通知的服务种类相应的缓冲器 44d 和 45d 由虚线指示。在中继节点，上述格式的 FR 帧 39 首先连接到每条数据链路专用的分类部件 43。分类部件 43 从该帧的地址域读数据链路连接标识符，并且选择与该标识符所指的虚拟通道相应的服务种类。虚拟通道和各服务种类存储在表 T 中。在已经完成的分类的基础上，分类部件 43 将每个帧加到与帧的服务种类相应的输入缓冲器中。每条输入数据链路因而具有四个输入缓冲器，每个服务种类 1 至 3 一个，并且拥挤通知有一个。每条数据链路具有一个专用的选择器 S1，它从服务种类专用的缓冲器中选择帧，并将它们送到节点内。在中继节点的输出

侧，将帧连接到与其自身数据链路相应的接口。在此接口，根据在节点输入侧所选择的服务种类，将它们提供到服务种类专用的输出缓冲器之一，从这里，选择器 S2 又将帧读到中继连接中。因而每输出数据链路有四个输出缓冲器，每个服务种类 1 至 3 一个并且拥挤通知有一个。另一方案是，可以为每条数据链路，甚至在节点的输出侧，单独设置分类部件，在此情形下，不必在节点内传送分类数据。

上述通信控制允许单个节点(最终是整个网络)的中继容量在不同用户之间得以合理划分。即使通信的突发性引起瞬时拥挤，本发明的方法也允许通过对发送突发串的通道所发送的另外通信进行缓冲而有效地控制通信。这样，在所承担的通信的限度内工作的通道仍使其通信经过短的延迟而通过网络。从而在每条虚拟通道上的通信围绕值 B_c 变化。

在连续拥挤时，过程是类似的，只是当虚拟通道专用的缓冲器溢出时必须放弃超过网络中继容量的通信。甚至在这样的情形下，放弃通信影响使网络过载的虚拟通道；其他通道的通信几乎根本不必减慢。为了放弃帧，使用在共同未决的芬兰专利申请 No.93xxxx 中所公开的方法是有利的，根据该方法，当在已满的缓冲器中接收一个帧时，放弃缓冲器的全部内容。为更详细地描述该方法，参阅上述共同未决专利申请。

如果在网络中有空闲中继容量，则不发送拥挤通知，从而通道带宽可以增至为其设置的最大值 $B_c + B_e$ 。在该情形下，承担的通信和超量的通信两者都被从缓冲器读到网络内部。从而即使在静态时刻，也将利用网络的容量，并且正确地处理另外的通信。

尽管已经参照附图中所示的例子描述了本发明，但是显然，本发明并不限于这些例子是明显的，而是可以在以上所公开的发明构思和后附权利要求的范围内进行修改。具体地说，例如，节点的内部结构可以以许多方式改变，即使根据本发明的思想进行基于拥挤通知的调节。例如，在间隔 $T_c/3$ 内，选择器 S3 能够对每个虚拟通道专用的缓冲器读一次或几次。拥挤通知的严重度也可以与其拥挤

等级受到连续监控的任何资源有关。在这种情形下，例如，能够采用在芬兰专利申请 No.925670 中所公开的方法，根据该方法，对网络的一个资源确定拥挤等级，拥挤等级的值与服务等级具有某种关系。

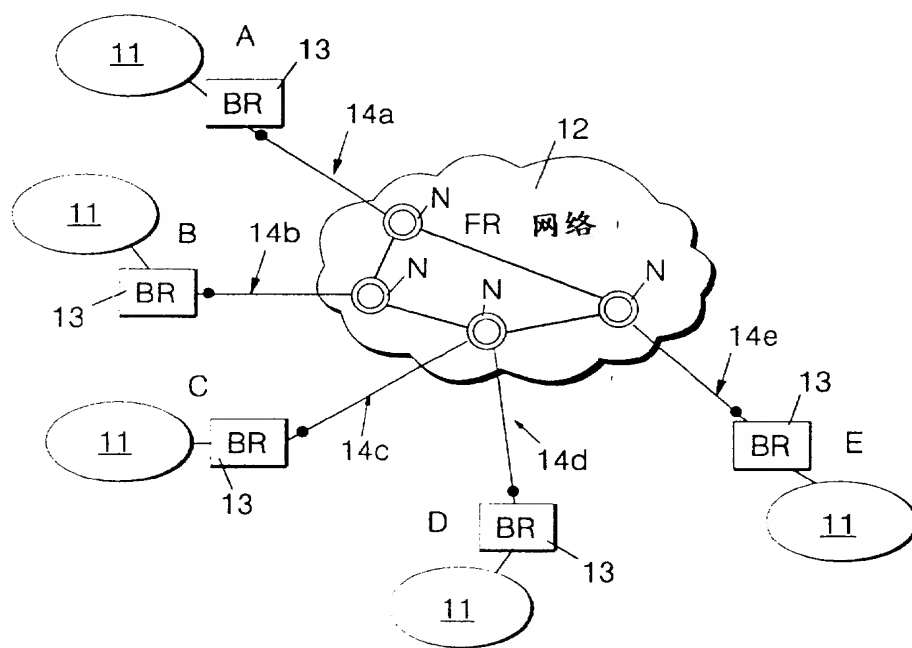


图. 1




图. 2




图. 3

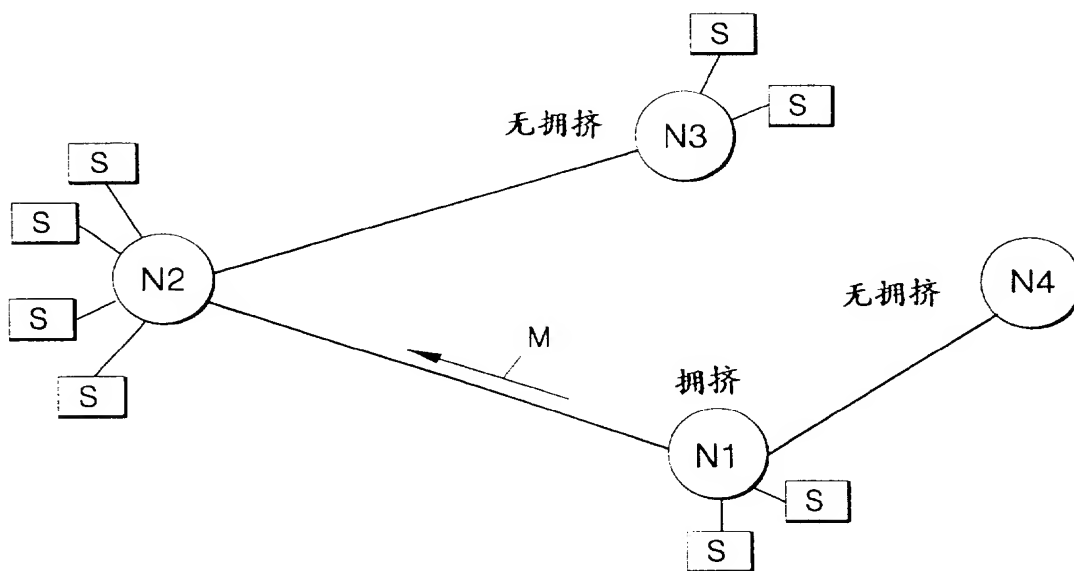





图. 4





US005970048A

United States Patent [19]**Pajuvirta et al.**[11] **Patent Number:** **5,970,048**[45] **Date of Patent:** **Oct. 19, 1999**[54] **CONTROL OF OVERLOAD SITUATIONS IN
FRAME RELAY NETWORK**[75] Inventors: **Juha Pajuvirta; Heikki Salovuori,**
both of Helsinki; **Jorma Matkaselkä ,**
Vantaa; **Mika Kasslin; Seppo**
Pyhälammii, both of Helsinki; **Mikko**
Olkkonen, Espoo; **Richard Fehlmann;**
Mikko Laiho, both of Helsinki, all of
Finland0 544 452 6/1993 European Pat. Off. .
0 576 122 12/1993 European Pat. Off. .
03034039 2/1991 Japan .
05103046 4/1993 Japan .
92/07438 4/1992 WIPO .
93/05596 3/1993 WIPO .
93/14605 7/1993 WIPO .
PCT/FI93/
00537 11/1993 WIPO .**OTHER PUBLICATIONS**[73] Assignee: **Nokia Telecommunications Oy,** Espoo,
Finland[21] Appl. No.: **08/647,955**[22] PCT Filed: **Nov. 29, 1994**[86] PCT No.: **PCT/FI94/00535**§ 371 Date: **Sep. 24, 1996**§ 102(e) Date: **Sep. 24, 1996**[87] PCT Pub. No.: **WO95/15637**PCT Pub. Date: **Jun. 8, 1995**[30] **Foreign Application Priority Data**

Nov. 30, 1993 [FI] Finland 935365

[51] Int. Cl.⁶ **H04J 3/14**[52] U.S. Cl. **370/230; 370/236; 370/466;**
370/412[58] **Field of Search** 370/229-236,
370/400, 412, 468, 420, 463; 340/825,
825.06, 825.07, 825.02[56] **References Cited****U.S. PATENT DOCUMENTS**5,115,429 5/1992 Hluchyj et al. 370/231
5,163,046 11/1992 Hahne et al. 370/237
5,436,891 7/1995 Grossman et al. 370/231
5,638,359 6/1997 Peltola et al. 370/229
5,694,390 12/1997 Yamato et al. 370/230**FOREIGN PATENT DOCUMENTS**


0 499 150 8/1992 European Pat. Off. .

A. Platt & M.J. Morse, Traffic Management in Frame Relay
Networks, Computer Networks and ISDN Systems, vol. 23,
No. 4 (Jan. 1992), pp. 305-316.Nobuyuki Tokura, et al., Congestion Avoidance Network,
Technical Report of the Institute of Electronics Information
and Communication Engineers, IN93-79 Oct. 1, 1993, pp.
49-54.Kuze, et al., A Study of Congestion Control in Frame Reply
Network, Proceedings of the 1993 IEICE Fall Conference
(Sep. 1993), vol. 3, B-388.Zenichi Yashiro, et al., Book for Explanation by Pictures of
Frame Relay, Ohm-sha, Aug. 30, 1993, pp. 49-63 and pp.
83-89.Official Action, dated Jul. 28, 1998, from the Japanese Patent
Office in related Application No. 7-515424.

(List continued on next page.)

Primary Examiner—Douglas W. Olms*Assistant Examiner*—Seema S. Rao*Attorney, Agent, or Firm*—Pillsbury Madison & Sutro LLP[57] **ABSTRACT**

A method and a system for congestion management in a frame relay network, which includes a) buffering of data at the input boundary of a subscriber node to virtual-channel-specific buffers, b) transmitting congestion notifications in a backward direction from the network nodes to the subscriber node of the moment, and c) controlling the amount of traffic supplied towards the network from the subscriber node buffer.

7 Claims, 3 Drawing Sheets

OTHER PUBLICATIONS

Jaakonmäki, Leena, "Lähiverkko-oppa".

CCITT, The International Telegraph and Telephone Consultative Committee, "Frame Mode Bearer Services", I.233, Geneva, 1992.

CCITT, The International Telegraph and Telephone Consultative Committee, "ISDN Data Link Layer Specification for Frame Mode Bearer Services", Q.922, Geneva, 1992.

CCITT, The International Telegraph and Telephone Consultative Committee, "Congestion Management for the ISDN Frame Relaying Bearer Service", I.370, Geneva, 1991.

"An Overview of Frame Relay Technology", *System Design*, Apr. 1991.

Grimshaw, Michael, "LAN Interconnections Technology" *Feature*, Jun. 1991, pp. 37-46.




FIG. 1




FIG. 2




FIG. 3

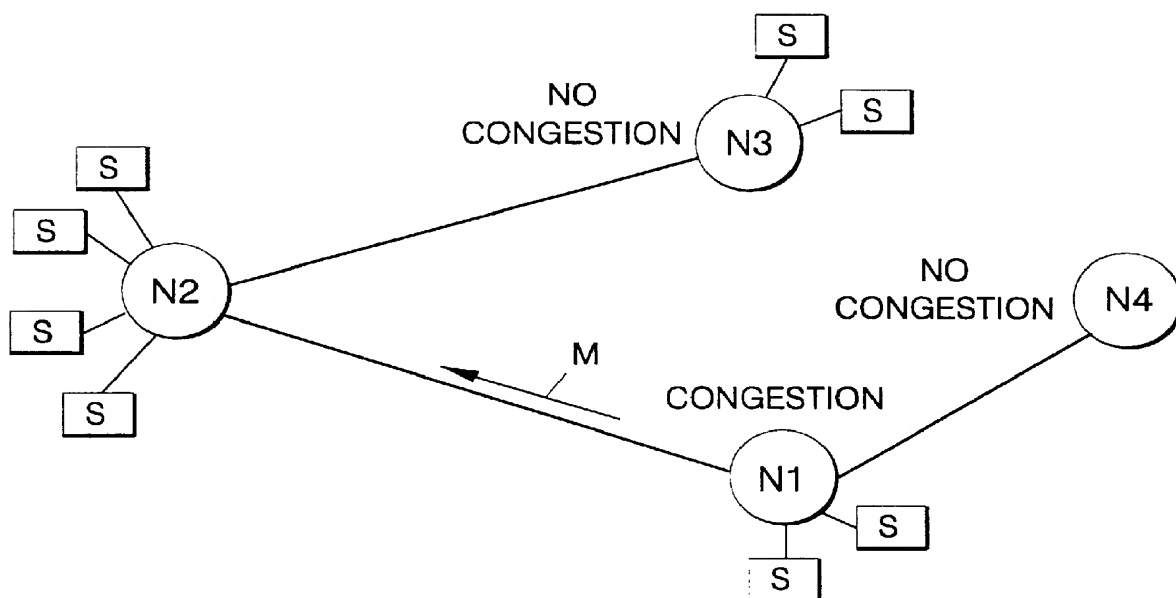





FIG. 4



CONTROL OF OVERLOAD SITUATIONS IN FRAME RELAY NETWORK

This application claims benefit of International application PCT/ FI94 /00535 filed Nov. 29, 1994.

BACKGROUND OF THE INVENTION

The invention relates to a method and a system for congestion management in a frame relay network.

Congestion means a situation in which the number of transmission requests exceeds the transmission capacity at a certain network point (called a bottle-neck resource) at a specific time. Congestion usually results in overload conditions, as a result of which the buffers overflow, for instance, and so packets will be retransmitted either by the network or the subscriber. The function of congestion management (CM) is to maintain a balance between the transmission requests and the transmission capacity so that the bottle-neck resources operate on an optimal level, and the subscribers are offered service in a way that assures fairness.

Congestion management can be divided into congestion avoidance (CA) and congestion recovery (CR). Congestion avoidance methods aim at preventing the occurrence of congestion in the network by dynamically adjusting the bandwidth of the subscribers in accordance with network congestion conditions and/or by altering the network routing so that part of the traffic load of the bottle-neck resources is shifted to idle resources. The purpose of recovery methods, in turn, is to restore the operation of the bottle-neck resources to an optimal level if the avoidance methods have failed to prevent the occurrence of congestion.

The frame relay (FR) technique is a packet-switched network technique used for the transmission of frames of varying length in place of the packet-switched network connections earlier placed in use. The protocol (X.25) applied generally in the present packet-switched networks requires plenty of processing and the transmission equipment is expensive, as a result of which the speeds also remain low. These matters are due to the fact that the X.25 standard was developed when the transmission connections used were still rather prone to transmission errors. The starting point of the frame relay technique was a considerably lower transmission line error probability. It has therefore been possible to discard a number of unnecessary functions in the frame relay technique, which makes the delivery of frames rapid and efficient. The Frame Mode Bearer Service is described generally in CCITT recommendation I.233 (Reference 1) and the associated protocol in recommendation Q.922 (Reference 2). Congestion in an FR network and congestion management mechanisms are described in CCITT recommendation I.370 (Reference 3). For a more detailed description of the FR technique, reference is made to *An Overview of Frame Relay Technology*, Datapro Management of Data Communications, McGraw-Hill Incorporated, April 1991, (Reference 4) as well as the above-mentioned recommendations.

In the FR network hierarchies presently in use, nodes have reception and transmission buffers corresponding only to the physical channels, i.e. the traffic of several different virtual channels and applications passes through the same buffer. Buffers are emptied as far as possible to links outbound from the node irrespective of the total level of congestion in the network. Links outbound from the node are thus loaded as much as possible even if the frame will probably be discarded closer to the middle of the network (at a congested node). In addition to wasting the network resources, dis-

carding of frames at trunk nodes of the network affects the applications using the network in the form of longer throughput delays. (The virtual channel refers to a virtual connection portion having the length of one transmission link while the virtual connection is the actual packet-switched end-to-end FR connection.)

Even though the traffic received from virtual channels is monitored on subscriber connections on the basis of general service parameters Bc, Be and CIR, the monitoring is not effected when frames are forwarded from the subscriber node. (Parameter Bc (committed burst size) represents the maximum amount of data the subscriber can transmit over the network within a certain time slot Tc (usually Tc=1s); parameter Be (excess burst size) represents the amount of data by which the subscriber can exceed the value Bc within the time slot Tc; and parameter CIR (committed information rate) represents the data transmission rate guaranteed by the network under normal conditions, CIR=Bc/Tc. These parameters are defined in Reference 3.) From the reception buffer, the frame is routed to the correct channel-specific transmission buffer. Frames arriving at the node thus pass through the entire node on the FIFO principle. Due to the burstiness of the traffic in a frame relay network, it often follows from the above that one virtual channel takes a major part of the buffering and relay capacity of a node. This affects the other virtual channels supplying traffic to the node both as longer throughput delays and as higher frame loss probabilities. Even one virtual channel with a highly bursty or unsatisfactory nature may cause the other virtual channels of the connection to be subjected to a considerable fall in the service level.

SUMMARY OF THE INVENTION

The object of the present invention is to obviate the drawbacks described above and to provide a new type of congestion management method in an FR network to allow the network resources to be utilized more fairly and more efficiently than before.

The idea of the invention is to transmit congestion notifications from the network nodes to the subscriber node of the virtual channel whose frame is received at the node, and to control the amount of traffic supplied by each virtual channel from a subscriber node towards the network in virtual-channel-specific buffers of the subscriber node, for instance according to the total level of congestion in the network indicated by these congestion notifications.

The method of the invention allows the relay capacity of a single node and, above all, of the entire network to be equitably divided between all subscribers.

BRIEF DESCRIPTION OF THE DRAWINGS

In the following, the invention and its preferred embodiments will be described in greater detail with reference to the examples illustrated in the accompanying drawings, in which

FIG. 1 illustrates a typical operating environment of the method according to the invention,

FIG. 2 illustrates an FR network subscriber node according to the invention,

FIG. 3 illustrates the format of a frame to be delivered in an FR network,

FIG. 4 illustrates the delivery of congestion notifications in the network,

FIG. 5a illustrates the changing of bandwidth according to congestion notifications in a first exemplary case,

FIG. 5b illustrates the changing of bandwidth according to congestion notifications in a second exemplary case, and

FIG. 6 illustrates an FR network trunk node according to the invention.

DETAILED DESCRIPTION

A frame relay network can be used by several different applications, which do not require similar services. It is therefore advantageous to employ the method of the invention in a network where the services are divided into different classes according to the applications, taking into account the two most significant parameters (frame loss probability and delay). Such a solution is disclosed in Finnish Patent Application No. 925671. In this application, it is suggested that the services be divided into three classes as follows:

- the first service class (class 1) offers interactive service where the delay is kept short,
- the second service class (class 2) offers a low frame loss probability without the delay having any great significance, and
- the third service class (class 3) offers both a short delay and a low frame loss probability.

Each trunk connection of a network realized in this way has service-class-specific buffers, one for each service class. A subscriber node, in turn, has virtual-channel-specific buffers on the side of the subscriber connection. These solutions will be described more closely below; otherwise reference is made to the Finnish patent application cited above.

FIG. 1 shows an FR network offering public network services, that is, a frame relay network 12 interconnecting local area networks 11 of different offices A . . . E of a single corporation or a plurality of corporations. The local area network 11 of each office has access to the FR service via a local area network bridge 13 and a data link indicated with the references 14a . . . 14e, respectively. The connection between an FR subscriber A . . . E and an FR network node N is known per se, wherefore it will not be described more closely herein. More detailed information about local area networks and bridges used in their interconnection can be found, e.g. in an article by Michael Grimshaw *LAN Interconnections Technology*, Telecommunications, February 1991, and in Lähiverkko-opas, Leena Jaakonmäki, Suomen ATK-kustannus Oy, 1991, which are incorporated herein by reference.

A typical feature of the known node structure of the FR network is that the same buffer is used for all frames, assuming that they are routed to the same physical connection. According to the present invention, on the contrary, buffers corresponding to the above-described service classes are provided at the output boundary of all network nodes and at the input boundary having trunk connections. FIG. 2 illustrates this kind of solution at a trunk node in the network. The node receives FR frames originally assembled in the bridges 13 of the subscriber connections (FIG. 1). The frame of the subscriber LAN 11 is inserted into the information field of the FR frame in the bridge 13 (with the exception of timing bits and other similar bits). FIG. 3 illustrates the insertion of a LAN frame 38 into the information field of an FR frame 39. It also shows a typical FR network frame format where the address field preceding the information field comprises two octets (bits 1 to 8). The bits 3 to 8 of the first octet and the bits 5 to 8 of the second octet form a data link connection identifier DLCI, which indicates to the node e.g. the virtual connection and virtual channel to which a particular frame belongs. The virtual channels are

distinguished from each other by means of the data link connection identifier. The data link connection identifier, however, is unambiguous only over a single virtual channel, and it may change in the node on transition to the next virtual channel. The bit 2 of the second address field octet, called a DE bit (Discard Eligibility Indicator), is also significant for the discarding of frames. In accordance with the CCITT recommendation, it is allowable to discard a frame, e.g. under congestion conditions if the DE bit of the frame has been set to one. As the other bits in the FR frame are not relevant to the present invention, they will not be described more closely herein. References 2 and 4 mentioned above are referred to for a more detailed description.

At a subscriber node on the periphery of the network (FIG. 2), the subscriber connections 14a, 14b, etc., (which in the example illustrated in FIG. 2 are connected to the same node) are connected at first to an identification unit 61, which receives FR frames formed in bridges 13 (FIG. 1). The identification unit 61 reads the data link connection identifier DLCI from the address field of the frame, and forwards the frame to an input buffer 62₁ . . . 62_n, corresponding to the virtual connection indicated by the identifier. Each data link has a specific selector S3 which selects frames from the input buffers of each virtual channel and forwards them to a centralized router 16, which routes the frames further to a classification unit 43 of the correct data link (the figure shows only one outbound data link). The classification unit 43 reads the identifier DLCI from the address field of the frame and selects from table T the service class corresponding to the virtual channel indicated by the identifier. On the basis of the classification it has completed, the classification unit 43 applies each frame to an output buffer 64a, 64b or 64c corresponding to the service class of the frame. Each outbound data link thus has three output buffers, one for each service class. The selector S4 reads the frames from these service-class-specific buffers further to the trunk connection.

Traffic transmitted by the subscribers over the FR network is thus buffered on the input side of the subscriber node specifically for each virtual connection. Incoming frames 39 are chained dynamically over each virtual connection. Depending on the service class of the virtual connection, the chain length has a predetermined allowable maximum value; this is smaller in service classes 1 and 3 and greater in service class 2. The selector S3 reads the buffers 62₁ . . . 62_n, e.g. in proportion to the amount of traffic allocated to them, whereby the fairness principle is met.

According to the invention, the amount of traffic supplied to the network 12 by each virtual channel is adjusted in the virtual-channel-specific buffers 62₁ . . . 62_n so that the amount of traffic varies depending on the level of congestion in the network around a basic value sold to the channel in question. The total amount of traffic supplied to the inner parts of the network is thus based on the service level sold to the virtual channel, on bandwidth, and on the total level of congestion in the network; of these, the service level and the bandwidth determine a certain basic value, around which the amount of traffic is adjusted according to what the total level of congestion is in the network. The adjustment is performed by transmitting congestion notifications from the network to the subscriber node of the virtual channel whose frame was received; the amount of traffic supplied towards the network from the subscriber node buffer corresponding to this virtual channel within a given time interval of a predetermined length is controlled according to the contents of the congestion notifications received from the network during this interval. The congestion notifications, indicated

in FIG. 2 by reference M, are connected to the control unit 65 controlling the selector S3 in the subscriber node, said control unit controlling the amount of data read from each virtual-channel-specific buffer according to the contents of the congestion notifications. Each buffer is thus emptied on the basis of its own parameters and its own congestion notifications.

As regards the subscriber node, the method of the invention can be compared to a floodgate: the traffic which each virtual-channel-specific buffer allows to be transmitted to the network is controlled (by means of the control unit 65) by a system which functions as a floodgate, restricting the amount of data read from the buffer, i.e. adjusting the bandwidth of the virtual channel.

This floodgate is a system which is bound to the service parameter Bc described above, and which allows only a certain amount of data to be transmitted from a virtual-channel-specific buffer towards the network at a certain period of time. As to the resolution of monitoring the bandwidth of the virtual channel, it is possible to use a time interval having the length of, e.g., $T_c/3$ (T_c has often the length of 1s). During each interval having the length of $T_c/3$, the buffer is emptied to an outbound connection towards the network by an amount allowed by the height of the floodgate. The rest of the frames are left in the buffer to wait for the next interval of $T_c/3$.

According to the invention, the nodes of the network transmit the above-mentioned congestion notifications M in a backward direction to the subscriber nodes. These congestion notifications indicate the fill rate of the buffers at the node in question. This mechanism is illustrated in FIG. 4. The nodes of the network are indicated by references N1 . . . N4, and the subscribers by references S. When a service-class-specific buffer at a network node, e.g. at node N1 in FIG. 2, exceeds a certain fill rate, e.g. 20%, the node transmits a congestion notification in a backward direction to the subscriber node (N2 in FIG. 4) of the virtual channel whose frame was received. The node (N1) identifies the correct virtual channel as described above by means of the data link connection identifier DLCI.

The height of the floodgate determining the emptying rate of each virtual-channel-specific buffer at the subscriber node (i.e. the bandwidth offered to the virtual channel) is adjusted according to the contents of the congestion notifications received from the network on said virtual channel in such a way that a higher load corresponds to a smaller bandwidth. Adjustment of the bandwidth during the interval $T_c/3$ is preferably performed according to the most severe congestion notification received during the respective interval. If the virtual channel does not receive any congestion notification during the interval $T_c/3$, the bandwidth of the channel is automatically increased according to the invention at the beginning of the following interval. The increase can be determined to be, e.g., 10%. The initial value for the bandwidth is the value $B_c + B_e$ sold to the channel, i.e. the maximum which the amount of traffic is not allowed to exceed under any circumstances.

By congestion notifications it is possible to indicate, for example, three different levels of congestion, whereby the changes in bandwidth corresponding to them (which are all calculated on the basis of the maximum value) may be, for instance, similar to those indicated in the following table.

Degree of severity of congestion notification	Change in bandwidth
1	-10%
2	-15%
3	-20%

To illustrate the principles described above, FIGS. 5a and 5b show two different examples of adjustment according to the invention. The vertical axis represents the height of the floodgate (i.e. bandwidth), whereas the horizontal axis represents time consisting of successive intervals having the length of $T_c/3$. In the case illustrated in FIG. 5a, a congestion notification M with severity degree 1 is at first received, wherefore the bandwidth is reduced immediately. After that a second congestion notification M with severity degree 2 is received, wherefore the bandwidth is again reduced immediately. The last, third notification M received during the interval $T_c/3$ has the degree of severity 3, wherefore the bandwidth is again reduced immediately (having then the lowest possible value). As no congestion notifications are received during the following interval $T_c/3$, the bandwidth of the channel grows automatically (instant T2) after this interval (at the beginning of the following interval $T_c/3$).

In the case illustrated in FIG. 5b, the order of the congestion notifications in time domain is reverse. At first a congestion notification with severity degree 3 is received, wherefore the bandwidth is reduced immediately to the lowest possible value. The other congestion notifications received during the interval do not change the bandwidth further, no matter what their degree of severity is. If a congestion notification with severity degree 2 were received at first, only a congestion notification with severity degree 3 would cause the bandwidth to be further reduced.

According to the invention, the bandwidth is thus reduced immediately upon reception of a congestion notification which is the most severe received so far during the interval in question, but the bandwidth can grow only if no congestion notifications are received during the interval $T_c/3$. In this case, the growth takes place after the interval in question has ended.

If the bandwidth has exceeded the value Bc, it drops according to a preferred embodiment of the invention to the value Bc immediately on reception of the first congestion notification, wherefore the abovedescribed limits set for the changes are not followed in this special case. This ensures rapid reaction above all to instantaneous congestion. In all cases, however, the bandwidth is adjusted upwards as described above.

Each degree of severity of a congestion notification, bound to the fill rate of the buffers located at the network nodes, can be set to correspond e.g. to the fill rate limits given in the following table:

Fill rate calculated on total buffer capacity	Degree of severity of congestion
20 . . . 39%	1
40 . . . 59%	2
60 . . . 100%	3

Each node of the network monitors the fill rates of the service-class-specific buffers continuously. When a new frame is received at a congested node, the node transmits a congestion notification M in the direction from which the frame was received, to the subscriber node of the virtual

channel in question. If the fill rate of the buffer is less than 20%, an incoming frame does not lead to the transmission of a congestion notification. Simultaneously with the transmission of a congestion notification, an interval **T1** of a predetermined length is set by a timer at the node; during this interval, no new congestion notifications are transmitted to the virtual channel in question. The length of the interval **T1** may be, for example, 100 ms (i.e. about one third of the interval $T_c/3$). In this way it is ensured that several congestion notifications are not transmitted in vain upon one burst occurring on the virtual channel. When the timer has timed off, it is again possible, if necessary, to transmit a congestion notification to the virtual channel.

Congestion notifications **M** should be delivered to egress nodes extremely fast to enable rapid reaction to congestion. According to a preferred embodiment of the invention, these congestion notifications form a separate, fourth service class, for which separate service-class-specific buffers are provided at the nodes. As regards the subscriber node, this means that the output side of the node is provided with a fourth output buffer, indicated in FIG. 2 by broken lines and the reference **64d**.

As regards the trunk node of the network, this embodiment is shown in FIG. 5, in which the buffers **44d** and **45d** corresponding to the service class of congestion notifications are indicated by broken lines. At the trunk node, an FR frame **39** of the format described above is connected at first to a classification unit **43** specific for each data link. The classification unit **43** reads the data link connection identifier from the address field of the frame, and selects the service class corresponding to the virtual channel indicated by the identifier. The virtual channels and the respective service classes are stored in a table **T**. On the basis of the classification it has completed, the classification unit **43** applies each frame to an input buffer corresponding to the service class of the frame. Each inbound data link has thus four input buffers, one for each service class 1 to 3 and one for congestion notifications. Each data link has a specific selector **S1** which selects the frames from the service-class-specific buffers and forwards them within the node. On the output side of the trunk node, the frames are connected to an interface corresponding to their own data link. At this interface they are supplied in accordance with the service class selected on the input side of the node to one of the service-class-specific output buffers, from which the selector **S2** reads the frames further to the trunk connection. Each outbound data link has thus four output buffers, one for each service class 1 to 3 and one for congestion notifications. Alternatively, classification units may be provided separately for each data link even on the output side of the node, in which case classification data need not be transferred within the node.

The above-described traffic control allows the relay capacity of a single node and, above all, of the entire network to be equitably divided between the different subscribers. In the event of instantaneous congestion, caused by burstiness of the traffic, the method of the invention allows the traffic to be efficiently controlled by buffering the additional traffic transmitted by the channel that sent the burst. Thus the channels operating within the limits of committed traffic still get their traffic through the network with a short delay. The traffic on each virtual channel thus varies around the value B_c .

In the event of continuous congestion, the procedure is otherwise similar but the traffic exceeding the relay capacity of the network must be discarded when there is overflow in

the virtual-channel-specific buffers. Even in such a case, the discarding of traffic affects the virtual channels which overload the network; the traffic of the other channels hardly needs to be slowed down at all. For discarding frames, it is advantageous to use the method disclosed in co-pending Finnish Patent Application No. 935364 (now Finnish patent 94814, corresponding to U.S. patent application Ser. No. 08/647, 950, filed Sep. 24, 1996 now pending), according to which the entire contents of a buffer are discarded when a frame is received in a buffer which is full. For a more detailed description of this method, reference is made to the above-mentioned co-pending patent application.

If there is idle relay capacity in the network, no congestion notifications are transmitted, whereby the channel bandwidths may grow to the maximum value $B_c + B_e$ set for them. In this case, both the committed traffic and the excess traffic are read from the buffers to the inner parts of the network. The capacity of the network will thus be utilized even in quiet times, and additional traffic will be treated correctly.

Although the invention has been described above with reference to the examples shown in the accompanying drawings, it will be obvious that the invention is not restricted to these examples, but can be modified within the scope of the inventive concept disclosed above and in the appended claims. In its details, the internal structure of the nodes, for example, may vary in many ways, even though the adjustment based on congestion notifications is performed according to the idea of the invention. During the interval $T_c/3$, for example, each virtual-channel-specific buffer can be read by the selector **S3** either only once or several times. The degree of severity of congestion notifications can also be bound to any resource whose congestion level is monitored continuously. In this case, it is possible to employ, for instance, the method disclosed in Finnish Patent Application No. 925670, according to which a congestion level is determined for a resource of the network, the value of the congestion level having a certain relationship with the values of the service levels.

We claim:

1. A method for congestion management in a frame relay network having a plurality of nodes between which data is transmitted, at least part of said nodes being subscriber nodes to which respective subscribers are connected over respective data links on respective virtual channels, said method comprising:

determining which of said virtual channels is associated with each frame to be transmitted, when the respective frame is received at a node, and

transmitting congestion notifications from respective ones of said nodes in a backward direction towards subscriber nodes to which respective subscribers are connected,

buffering data received from respective ones of said subscribers at a respective input boundary of each respective subscriber node, to respective virtual-channel-specific buffers,

said transmitting including transmitting respective ones of said congestion notifications each from a respective node to the respective said subscriber node of the respective said virtual channel a frame associated with which is received at that particular moment at the respective said node, and

controlling the amount of traffic supplied towards the network from each respective said subscriber node buffer corresponding to each respective said virtual channel during a respective certain interval having a

9

predetermined length on the basis of the contents of respective ones of said congestion notifications received from the network during each respective said interval, whereby traffic from respective ones of said subscribers to respective ones of said buffers is discarded, when necessary for managing congestion. 5

2. The method according to claim 1, further comprising:

binding the degree of severity of congestion notifications to the fill rate of respective ones of said buffers at a respective said node, whereby a certain severity degree corresponds to each fill rate range, and a change of constant magnitude in the amount of traffic supplied towards the network during a respective said interval corresponds to a certain severity degree of congestion. 10

3. The method according to claim 1, wherein: 15

if no congestion notifications belonging to a respective said virtual channel are transmitted during a respective said interval, increasing the amount of traffic at a respective said subscriber node by a certain constant value, without exceeding an allowed maximum value. 20

4. The method according to claim 3, wherein:

when a respective said amount of traffic exceeds a respective committed burst size, reducing the respective amount of traffic to a value corresponding to the respective committed burst size immediately upon reception of a respective said congestion notification. 25

5. The method according to claim 1, wherein:

said controlling includes controlling the respective said amount of traffic according to the most severe respective congestion notification received during a respective said interval. 30

10

6. The method according to claim 1, further including:

after the transmission a respective said congestion notification to a given said virtual channel, preventing transmission of a following congestion notification to the same given virtual channel for a predetermined time, in order to prevent several congestion notifications from being sent in vain upon one burst occurring on the respective given virtual channel.

7. A system for congestion management in a frame relay networks comprising:

a plurality of nodes arranged for transmission of data therebetween, at least part of said nodes being subscriber nodes to which respective subscribers of the network are connected over respective data links on respective virtual channels, the nodes comprising means for transmitting congestion notifications in a backward direction towards respective ones of the subscribers,

virtual-channel-specific buffers to which subscriber-generated data is buffered are provided at an input boundary of each subscriber node, and

each subscriber node comprises means for controlling the amount of subscriber-generated traffic supplied towards the network from a virtual-channel-specific buffer during a certain interval having a predetermined length, said controlling means being responsive to the contents of congestion notifications received from the network at the respective said subscriber node during said certain interval.

* * * * *